

Certification of Compact Low-Stretch Routing Schemes

Alkida Balliu^{*1} and Pierre Fraigniaud^{†2}

¹ Institut de Recherche en Informatique Fondamentale, CNRS and University Paris Diderot, France.

² Gran Sasso Science Institute, L'Aquila, Italy.

Abstract

On the one hand, the correctness of routing protocols in networks is an issue of utmost importance for guaranteeing the delivery of messages from any source to any target. On the other hand, a large collection of *routing schemes* have been proposed during the last two decades, with the objective of transmitting messages along short routes, while keeping the routing tables small. Regrettably, all these schemes share the property that an adversary may modify the content of the routing tables with the objective of, e.g., blocking the delivery of messages between some pairs of nodes, without being detected by any node.

In this paper, we present a simple *certification* mechanism which enables the nodes to locally detect any alteration of their routing tables. In particular, we show how to locally verify the stretch-3 routing scheme by Thorup and Zwick [SPAA 2001] by adding certificates of $\tilde{O}(\sqrt{n})$ bits at each node in n -node networks, that is, by keeping the memory size of the same order of magnitude as the original routing tables. We also propose a new *name-independent* routing scheme using routing tables of size $\tilde{O}(\sqrt{n})$ bits. This new routing scheme can be locally verified using certificates on $\tilde{O}(\sqrt{n})$ bits. Its stretch is 3 if using handshaking, and 5 otherwise.

1 Introduction

1.1 Context

A *routing scheme* is a mechanism enabling to deliver messages from any source to any target in a network. The latter is typically modeled as an undirected connected weighted graph $G = (V, E)$ where V models the set of routers and E models the set of communication links between routers. All edges incident to a degree- d node are labeled from 1 to d , in an arbitrary manner, and the label at a node u of an incident edge e is called the *port number* of edge e at u . A routing scheme consists of a way of assigning a routing *table* to every node of the given network. These tables should contain enough information so that, for every target node t , each node is able to compute the port number of the incident edge through which it should forward a message of destination t . The routing tables must collectively guarantee that every message of any source s and any target t will eventually be delivered to t .

Two scenarios are generally considered in the literature. One scenario allows the routing scheme to assign *names* to the nodes, and each target is then identified by its given name. The

^{*}Additional support from ANR Project DESCARTES. This work was done during the first author visit to Institut de Recherche en Informatique Fondamentale (IRIF) in 2016-2017.

[†]Additional support from ANR Project DESCARTES and Inria Project GANG.

other, called *name independent*, is assuming that fixed names are given a priori (typically, a name is restricted to be the identity of a node), and the scheme cannot take benefit of naming nodes for facilitating routing.

Among many criteria for evaluating the quality of routing schemes, including, e.g., the time complexity for constructing the routing tables, the two main parameters characterizing a routing scheme are the *size* of its routing tables, and the *stretch*. The stretch of a routing scheme is the maximum, taken over all pairs of source-target nodes, of the ratio between the length of the route generated by the scheme from the source to the target, and the length of a shortest path between these two nodes. During the last two decades, there has been an enormous effort to design *compact* routing scheme (i.e., schemes using small tables) of *low* stretch (i.e., with stretch upper bounded by a constant) – see, e.g., [1, 2, 3, 4, 8, 13, 17, 18, 19]. A breakthrough was achieved in [19] where almost tight tradeoffs between size and stretch were explicitly demonstrated. In particular, [19] showed how to design a routing scheme with tables of size $\tilde{O}(\sqrt{n})$ bits and stretch 3, in any network¹.

All the aforementioned routing schemes share the property that nodes do not have the capability to realize that the routing tables have been modified (either involuntarily or by an attacker). That is, a group of nodes may be provided with routing tables which look consistent with a desired routing scheme, but which do not achieve the desired performances of that scheme (e.g., large stretch, presence of loops, etc.). Indeed, the nodes are not provided with sufficient information to detect such an issue locally, that is, by having each node inspecting only the network structure and the tables assigned to nodes in its vicinity.

1.2 Objective

The objective of this paper is, given a routing scheme, to design a mechanism enabling each node to locally detect the presence of falsified routing tables, in the following sense. If one or more tables are erroneous, then at least one node must be able to detect that error by running a verification algorithm exchanging messages only between neighboring nodes.

Our mechanism for locally verifying the correctness of routing tables is inspired from proof-labeling schemes [16]. It is indeed based on assigning to each node a *certificate*, together with its routing table, and designing a distributed verification algorithm that checks the consistency of these certificates and tables by having each node inspecting only its certificate and its routing table, and the certificate and routing table of each of its neighbors. The set of certificates assigned to the nodes and the verification algorithm running at all nodes in parallel must satisfy that: (1) if all tables are correctly set, then, with some appropriate certificates, all nodes *accept*, and (2) if one or more tables are incorrectly set, then, for every assignment of the certificates, at least one node must *reject*. The second condition guarantees that the verification algorithm cannot be cheated: if the tables are incorrect, there are no ways of assigning the certificates such that all nodes accept.

Rephrasing the objective of the paper, our goal is to assign certificates to nodes, of size not exceeding the size of the routing tables, enabling the nodes to collectively verify the correctness of the routing tables, by having each node interacting with its neighbors only.

1.3 Our Results

We show how to locally verify the stretch-3 size- $\tilde{O}(\sqrt{n})$ routing scheme by Thorup and Zwick [19]. Our certification mechanism uses certificates of $\tilde{O}(\sqrt{n})$ bits at each node, that is, these certificates have size of the same order of magnitude as the original routing tables. Hence, verifying

¹The notations \tilde{O} and $\tilde{\Omega}$ ignore polylogarithmic factors.

the scheme in [19] can be done without modifying the scheme, and without increasing the memory space consumed by that scheme. We also show that the same holds for the whole hierarchy of routing schemes proposed in [19] for providing a tradeoff between size and stretch.

The situation appears to be radically different for name-independent routing schemes. The stretch-3 name-independent routing scheme by Abraham et al. [2] also uses tables of size $\tilde{O}(\sqrt{n})$ bits. However, each table includes references to far away nodes, whose validity does not appear to be locally verifiable using certificates of reasonable size. On the other hand, a simplified version of the scheme in [2] can be verified locally with certificates of size $\tilde{O}(\sqrt{n})$ bits, but its stretch becomes at least 7. Therefore, we propose a new name-independent routing scheme, with tables of size $\tilde{O}(\sqrt{n})$ bits that can be verified using certificates on $\tilde{O}(\sqrt{n})$ bits as well. This new routing scheme has stretch at most 5, and the stretch can even be reduced to 3 using handshaking². The routing scheme of Arias et al. [3] has also stretch 5, but it does not appear to be locally verifiable with certificates of reasonable size, and using handshaking does not enable to reduce the stretch.

All our results are summarized in Table 1.

scheme	stretch	name independent	verifiable	comment
[19]	3	no	yes	–
[3]	5	yes	?	–
[2]	3	yes	?	–
this paper	5	yes	yes	–
this paper	3	yes	yes	handshaking

Table 1: Summary of our results compared to previous work. All the listed routing schemes have space complexity of $\tilde{O}(\sqrt{n})$ bits. Our verification algorithms for these schemes use certificates on $\tilde{O}(\sqrt{n})$ bits.

1.4 Related Work

The design of *compact* routing tables, and the explicit identification of tradeoffs between the table size and the routes length was initiated thirty years ago, with the seminal work in [18] and [17]. Since then, a large amount of papers were published on this topic, aiming at refining these tradeoffs, and at improving different aspects of the routing schemes, including routing in specific classes of graphs (see [12, 14]). In particular, routing schemes were designed for trees in [8, 19], with space complexity $O(\log^2 n / \log \log n)$ bits³. This space complexity was shown to be optimal in [9].

It was proved [15] that, in n -node networks, any shortest path routing scheme requires tables of size $\tilde{\Omega}(n)$ bits. The aforementioned routing scheme in [19] with stretch 3 and space complexity $\tilde{O}(\sqrt{n})$ bits was shown to be optimal in [13], in the following sense: no routing scheme with space complexity $o(n)$ bits can achieve a stretch $s < 3$, and, assuming the correctness of a conjecture by Erdős regarding a tradeoff between girth and edge density in graphs, every routing scheme with stretch $s < 5$ has space complexity $\Omega(\sqrt{n})$ bits. On the positive side, [19] tightens the size-stretch tradeoff of [17] by showing that, for every $k \geq 2$, there exists a routing scheme with stretch $4k - 5$ and space complexity $\tilde{O}(n^{1/k})$ bits. (The stretch can be reduced to $2k - 1$

²The handshaking mechanism is similar to DNS lookup in TCP/IP. It allows querying some node(s) for getting additional information about the route to the target.

³The space complexity can be reduced to $O(\log n)$ if the designer of the routing scheme is also allowed to assign the port numbers to the nodes.

using handshaking). Recently, [6] showed that, for $k \geq 4$, a stretch $s = \alpha k$ with $\alpha < 4$ can be achieved using routing tables of size $\tilde{O}(n^{1/k})$.

The distinction between name-independent routing schemes, and routing schemes assigning specific names to the nodes was first made in [4]. Then, [5] presented techniques for designing name-independent routing schemes with constant stretch and space complexity $o(n)$ bits. Almost 15 years after, [3] described a name-independent routing scheme with stretch 5 and space complexity $\tilde{O}(\sqrt{n})$ bits. This was further improved in [2] thanks to a name-independent routing scheme with stretch 3 and space complexity $\tilde{O}(\sqrt{n})$ bits. A couple of years later, [1] showed that there are tradeoffs between stretch and space complexity for name-independent routing schemes as well. Specifically, [1] showed that, for any $k \geq 1$, there exists a name-independent routing scheme with space complexity $\tilde{O}(n^{1/k})$ bits and stretch $O(k)$.

The certification mechanism used in this paper is based on the notion of proof-labeling scheme introduced in [16] in which an oracle, called *prover*, assigns certificates to the nodes, and a distributed algorithm, called *verifier*, checks that this certificates collectively form a proof that the global state of the network is legal with respect to a given boolean network predicate. Variants of proof-labeling scheme have been considered in, e.g., [10] and [11]. More generally, see [7] for a survey of distributed decision.

1.5 Structure of the Paper

In Section 2, we recall the main concepts related to the design of routing schemes, and those related to distributed verification. Section 3 is dedicated to routing schemes assigning names to nodes: we briefly recall the main features of the scheme in [19], and then show how to locally verify that scheme using small certificates. In Section 4, we show that the name-independent routing scheme in [2] is unlikely to be locally verifiable with small certificates. Therefore, Section 5 presents a new name-independent routing scheme, that is shown to be locally verifiable with small certificates. This new scheme has space-complexity $\tilde{O}(\sqrt{n})$ bits, but stretch 5. Section 6 shows that the stretch-5 scheme in [3] is, like the one in [2], unlikely to be locally verifiable with small certificates. Finally, Section 8 concludes the paper with pointers to research directions worth to be investigated.

2 Definitions

2.1 Routing Schemes

Let \mathcal{F} be a family of edge-weighted graphs with edges labeled at each node by distinct port numbers from 1 to the degree of the node. The weights are all positive, and the weight of edge e represents its length. It is thus denoted by $\text{length}(e)$. The nodes are given distinct identities stored on $O(\log n)$ bits. For the sake of simplifying notations, we do not make a distinction between a node v and its identity, also denoted by v . Given two nodes u, v , we denote by

$$\delta(u, v) = \text{weighted distance between } u \text{ and } v.$$

Given an edge e of extremity u , the port number of e at u is denoted by $\text{port}_u(e)$.

A *routing scheme* for \mathcal{F} is a mechanism assigning a *name*, $\text{name}(u)$, and a routing *table*, $\text{table}(u)$, to every node u of every graph $G \in \mathcal{F}$ such that, for any pair (s, t) of nodes of any $G \in \mathcal{F}$, there exists a path u_0, u_1, \dots, u_k from s to t in G with $u_0 = s$, $u_k = t$, and

$$\text{table}(u_i)(\text{name}(t)) = \text{port}_{u_i}(\{u_i, u_{i+1}\}) \quad (1)$$

for every $i = 0, \dots, k-1$. That is, every intermediate node u_i , $0 \leq i < k$, can determine on which of its ports the message has to be forwarded, based solely on its routing table, and on

the name of the target. In Eq. 1, each table is viewed as a function taking names as arguments, and returning port numbers. The path u_0, u_1, \dots, u_k is then called the *route* generated by the scheme from s to t . It is worth pointing out the following observations.

- *Name-independent* routing schemes are restricted to use names that are fixed a priori, that is, the name of a node is its identity, i.e., $\text{name}(v) = v$ for every node v . Instead, name-dependent routing schemes allow names to be set for facilitating routing, and names are typically just bounded to be storable on a polylogarithmic number of bits.
- The *header* of a message is the part of that message containing all information enabling its routing throughout the network. The header of a message with destination t is typically $\text{name}(t)$. However, some routing schemes ask for message headers that can be modified. This holds for both name-dependent schemes (like, e.g., [19]) and name-independent schemes (like, e.g., [2]). The typical scenario requiring modifiable headers is when a message is routed from source s to target t as follows. From $\text{name}(t)$ and $\text{table}(s)$, node s can derive the existence of some node v containing additional information about how to reach t . Then the message is first routed from s to v , and then from v to t . Distinguishing these two distinct parts of the routes from s to t often requires to use different headers. In case of modifiable headers, Eq. (1) should be tuned accordingly as the argument of routing is not necessarily just a name, but a header.
- Some routing schemes may use a mechanism called *handshaking*, which is an abstraction of mechanisms such as Domain Name System (DNS) enabling to recover an IP address from its domain name. Let us consider the aforementioned scenario where a routing scheme routes a message from s to t via an intermediate node v identified by s from $\text{name}(t)$ and $\text{table}(s)$. One can then enhance the routing scheme by a handshaking mechanism, enabling s to query v directly, and to recover the information stored at v about t at no cost. Then s can route the message directly from s to t , avoiding the detour to v . Handshaking is used in [19] to reduce the stretch of routing schemes with space complexity $\tilde{O}(n^{1/k})$ bits from $4k - 5$ to $2k - 1$, for every $k > 2$.

The *size* of a routing scheme is the maximum, taken over all nodes u of all graphs in \mathcal{F} , of the memory space required to encode the function $\text{table}(u)$ at node u .

The *stretch* of a routing scheme is the maximum, taken over all pairs (s, t) of nodes in all graphs $G \in \mathcal{F}$, of the ratio of the length of the route from s to t (i.e., the sum of the edge weights along that route) with the weighted distance between s and t in G .

2.2 Distributed Verification

Given a graph G , a *certificate function* for G is a function

$$\text{certificate} : V(G) \rightarrow \{0, 1\}^*$$

assigning a certificate, $\text{certificate}(u)$, to every node $u \in V(G)$. A *verification* algorithm is a distributed algorithm running concurrently at all nodes in parallel. At every node $u \in V(G)$ of every graph $G \in \mathcal{F}$, the algorithm takes as input $\text{certificate}(u)$ and $\text{table}(u)$ assigned to node u , as well as the collection of pairs $(\text{table}(v), \text{certificate}(v))$ assigned to all neighbors v of u , and outputs *accept* or *reject*.

Definition 1 A routing scheme for \mathcal{F} is verifiable if there exists a verification algorithm VERIF such that, for every $G \in \mathcal{F}$,

- if the tables given to all nodes of G are the ones specified by the routing scheme, then there exists a certificate function for G such that the verification algorithm VERIF outputs *accept* at all nodes;

- if some tables given to some nodes of G differ from the ones specified by the routing scheme, then, for every certificate function for G , the verification algorithm `VERIF` outputs reject in at least one node.

The second bullet guarantees that if an adversary modifies some routing tables, or even just a single bit of a single table, then there are no ways it can also modify some, or even all certificates so that to force all nodes to accept: at least one node will detect the change. Of course, this node does not need to be the same for different modifications of the routing tables, or for different certificates.

3 Name-dependent routing scheme

In this section, we show how to verify the stretch-3 routing scheme by Thorup and Zwick in [19]. This scheme uses tables of $\tilde{O}(\sqrt{n})$ bits of memory at each node. We show the following:

Theorem 2 *The stretch-3 routing scheme by Thorup and Zwick in [19] can be locally verified using certificates of size $\tilde{O}(\sqrt{n})$ bits.*

Before proving the theorem, we first recall the general structure of the routing scheme in [19].

3.1 The routing scheme

The routing scheme assigns names and tables to the nodes of every $G = (V, E)$ as follows.

3.1.1 Landmarks, Bunch and Clusters

The routing scheme in [19] uses the notion of *landmarks* (a.k.a. *centers* in the “older” terminology of [19]). These landmarks form a subset $L \subseteq V$ of nodes. For $v \in V$, let l_v denote the landmark closest to v in G . For every $v \in V$, the *bunch* of v with respect to the set L is defined as follows.

$$\text{bunch}(v) = \{u \in V : \delta(v, u) < \delta(l_v, v)\}.$$

The routing scheme in [19] also uses the notion of *cluster*. For every node $v \in V$,

$$\text{cluster}(v) = \{u \in V : \delta(v, u) < \delta(l_u, u)\}.$$

As a consequence, for every $u, v \in V$, we have

$$u \in \text{cluster}(v) \iff v \in \text{bunch}(u).$$

Note that since, for every $v \in V$, and every $l \in L$, we have $l \notin \text{bunch}(v)$, it follows that $\text{cluster}(l) = \emptyset$ for every $l \in L$. By construction of the bunches and the clusters, it also holds that, for every $v \in V$, $\text{cluster}(v) \cap L = \emptyset$. Also, clusters satisfy the following property.

Lemma 3 ([19]) *If $u \in \text{cluster}(v)$ then, for every node w on a shortest path between u and v , we have $u \in \text{cluster}(w)$*

In [19], the landmarks are chosen by an algorithm that samples them uniformly at random in V until the following holds: for every node v , $|\text{cluster}(v)| < 4\sqrt{n}$. It is proved that this algorithm returns, w.h.p., a set of landmarks of size at most $2\log(n)\sqrt{n}$.

3.1.2 Names and tables

For every two nodes v and t , let

$$\text{next}(v, t)$$

be the port number of an edge incident to v on a shortest path between v and t . Each node $t \in V$ is assigned a $3\lceil\log(n)\rceil$ -bit name as follows:

$$\text{name}(t) = (t, l_t, \text{next}(l_t, t)).$$

Each node $v \in V$ then stores the following information in its routing table, $\text{table}(v)$:

- the identities of all the landmarks $l \in L$;
- the identities of all nodes $t \in \text{cluster}(v)$;
- the set $\{\text{next}(v, t) : t \in L \cup \text{cluster}(v)\}$.

All these information can be stored using $\tilde{O}(\sqrt{n})$ bits.

3.1.3 Routing

The stretch-3 routing scheme then works as follows. Note that, by Lemma 3, any message from a node v to a node in $\text{cluster}(v)$ reaches its target along a shortest path. The same holds for landmarks since every node is given information about how to reach every landmark. In general, let us assume that v wants to send a message to some node t with label $(t, l_t, \text{next}(l_t, t))$ that is neither a landmark nor belongs to $\text{cluster}(v)$. In this case, v extracts the landmark l_t nearest to t from t 's name (note here the impact of allowing the scheme to assign specific names to nodes), and forwards the message through the port on a shortest path towards l_t using the information v has in its table. Upon reception of the message, l_t forwards the message towards t on a shortest path using $\text{next}(l_t, t)$ (this information can also be extracted from the name of t), to reach a node $z \in \text{bunch}(t)$. At last, since $z \in \text{bunch}(t)$, we have $t \in \text{cluster}(z)$, which means that z can route to t via a shortest path using the information available in its table. By Lemma 3, this also holds for every node along a shortest path between z and t . Using symmetry, and triangle inequality, [19] shows that this routing scheme guarantees stretch 3.

3.2 Proof of Theorem 2

In order to enable local verification of the stretch-3 routing scheme in [19], a certificate of size $\tilde{O}(\sqrt{n})$ bits is given to each node. Let $G = (V, E)$ be an undirected graph with positive weights assigned to its edges, and a correct assignment of the routing tables to the nodes according to the specifications of [19] as summarized in Section 3.1. Then each node $v \in V$ is assigned a certificate composed of:

- the distance between v and every landmark in L ;
- the distance between v and every node in $\text{cluster}(v)$;
- the set $\{\delta(t, l_t) : t \in \text{cluster}(v)\}$.

As claimed, all these information can be stored using $\tilde{O}(\sqrt{n})$ bits of memory.

We assume, without loss of generality, that all nodes know n (verifying the value of n is easy using a proof-labeling scheme with $O(\log n)$ -bit certificates [16]). The verification of the routing scheme then proceeds as follows. We describe the verification algorithm `VERIF` running

at node $v \in V$. This verification goes in a sequence of steps. At each step, either v outputs *reject* and stops, or it goes to the next step.

We denote by $L(v)$, $C(v)$, and $\{N(v, t) : t \in L(v) \cup C(v)\}$ the content of the routing table of v . These entries are supposed to be the set of landmarks, the cluster of v , and the set of next-pointers given to v , respectively. We also denote by d the distance given in the certificates. That is, node v is given a set $\{d(v, t) : t \in L(v) \cup C(v)\}$ and a set $\{d(t, l_t) : t \in C(v)\}$ where l_t is supposed to be the node in $L(v)$ closest to $t \in C(v)$. Of course, if a node does not have a table and a certificate of these forms, then it outputs *reject*. So, we assume that all tables and certificates have the above formats. The algorithm proceeds as follows at every node v . Node v checks that

1. the information in its table satisfy $|C(v)| \leq 4\sqrt{n}$ and $|L(v)| \leq 2 \log n \sqrt{n}$ bits;
2. it has the same set of landmarks as its neighbors;
3. for every $l \in L(v)$, there exists a neighbor u of v satisfying

$$d(v, l) = \text{length}(\{v, u\}) + d(u, l),$$

and all other neighbors satisfy

$$d(v, l) \leq \text{length}(\{v, u\}) + d(u, l),$$

with $N(v, l)$ pointing to a neighbor u satisfying $d(v, l) = \text{length}(\{v, u\}) + d(u, l)$;

4. if $v \in L(v)$ then $C(v) = \emptyset$;
5. $v \in C(v)$ and, for every node $t \in C(v)$, there exists a neighbor u of v satisfying $t \in C(u)$ with

$$d(v, t) = \text{length}(\{v, u\}) + d(u, t),$$

and every neighbor u of v with $t \in C(u)$ satisfies

$$d(v, t) \leq \text{length}(\{v, u\}) + d(u, t),$$

with $N(v, t)$ pointing to a neighbor u satisfying $t \in C(u)$ and $d(v, t) = \text{length}(\{v, u\}) + d(u, t)$;

6. for every node $t \in C(v)$, for every neighbor u of v satisfying $t \in C(u)$, the distance $d(t, l_t)$ in the certificate of u is equal to the distance $d(t, l_t)$ in the certificate of v ;
7. for every $t \in C(v)$, it holds that $d(v, t) < d(t, l_t)$;
8. for every neighbor u , and every $t \in \text{cluster}(u) \setminus \text{cluster}(v)$, it holds that

$$\text{length}\{v, u\} + d(u, t) \geq d(t, l_t).$$

If v passes all the above tests, then v outputs *accept*, else it outputs *reject*.

We now establish the correctness of this local verification algorithm, that is, we show that it satisfies the specification stated in Section 2.2. First, by construction, if all tables are set according to [19], that is, if, for every node v , $L(v) = L$, $C(v) = \text{cluster}(v)$ and $N(v, t) = \text{next}(v, t)$ for all $t \in L \cup \text{cluster}(v)$, then every node running the verification algorithm with the appropriate certificate $\{\delta(v, t) : t \in L \cup \text{cluster}(v)\} \cup \{\delta(t, l_t) : t \in \text{cluster}(v)\}$ where l_t is the node

in L closest to $t \in \text{cluster}(v)$, will face no inconsistencies with its neighbors, i.e., all the above tests are passed, leading every node to accept, as desired.

So, it remains to show that if some tables are not what they should be according to [19], then, no matter the certificates assigned to the nodes, at least one node will fail one of the tests.

If all nodes output *accept*, then, by Step 1, all routing tables are of the appropriate size. Also, by Step 2, the set L of landmarks given to the nodes is the same for all nodes, as otherwise there will be two neighbors that would have different sets. Moreover, by Step 3, we get that, at every node v , the distances of this node to the landmarks, as stored in its certificate, are correct, from which we infer that $N(v, l)$ is appropriately set in the table of v , that is $N(v, l) = \text{next}(v, l)$ for every $l \in L$. Hence, if all the tests in Steps 1-3 are passed, all the data referring to L in both the tables and the certificates are consistent. In particular, every node v knows the landmark l_v which is closest to it, and its distance $\delta(v, l_v)$.

We now show that if all these tests as well as the remaining tests are passed, then the clusters in the tables are correct, w.r.t. L , as well as the next-pointers. We first show that, if all tests are passed, then, for every node $v \in V$,

$$C(v) \subseteq \text{cluster}(v).$$

By Step 4, this latter equality holds for every landmark v . By Step 5, we get that, at every node v , the distance of this node to every node $t \in C(u)$, as stored in its certificate, are correct, from which we infer that $N(v, t)$ is appropriately set in the table of v , that is $N(v, t) = \text{next}(v, t)$ for every $t \in C(v)$. By Step 6, we get that, for every $t \in C(v)$, we do have $d(l_t, t) = \delta(l_t, t)$. Indeed, this equality will be checked by all nodes on a shortest path between v and t (whose existence is guaranteed by Step 5), and t has the right distance $\delta(t, l_t)$ in its certificate by Step 3. Recall that

$$\text{cluster}(v) = \{t \in V \mid \delta(v, t) < \delta(l_t, t)\}$$

where l_t is the landmark closest to t . Step 7 precisely checks that inequality.

It remains to show that there are no nodes in $\text{cluster}(v)$ that are not in $C(v)$. Assume that there exists $t \in \text{cluster}(v) \setminus C(v)$, and let P be a shortest path between v and t . Let v' be the closest node to v on P such that $t \in C(v') \subseteq \text{cluster}(v')$. Note that such a node v' exists as $t \in C(t)$. Let v'' be the node just before v' on P traversed from v to t . By Lemma 3, since $t \in \text{cluster}(v)$, we also have $t \in \text{cluster}(v'')$. We have $t \in \text{cluster}(v'') \setminus C(v'')$. Therefore $\delta(v'', t) < \delta(l_t, t)$. Now,

$$\delta(v'', t) = \text{length}(\{(v'', v')\}) + \delta(v', t)$$

because P is a shortest path between v'' and t passing through v' . So,

$$\text{length}(\{(v'', v')\}) + \delta(v', t) < \delta(l_t, t).$$

Step 8 guarantees that it is not the case. Therefore, there are no nodes in $\text{cluster}(v) \setminus C(v)$. It follows that $\text{cluster}(v) = C(v)$ for all nodes v . This completes the proof of Theorem 2.

4 Name-independent Routing Scheme

The purpose of this section is twofold. First, it serves as recalling basic notions that will be helpful for the design of our new name-independent routing scheme. Second, it is used to show why the known name-independent routing scheme in [2] appears to be difficult, and perhaps even impossible to verify locally.

4.1 The Routing Scheme of [2]

The stretch-3 name-independent routing scheme of [2] uses $\tilde{O}(\sqrt{n})$ space at each node. We provide a high level description of that scheme. Recall that, in name-independent routing, a target node is referred only by its identity. That is, $\text{name}(t)$ is the identity of t , i.e., $\text{name}(t) = t$.

Let $G = (V, E)$. For every node $v \in V$, the *vicinity ball* of v , denoted by $\text{ball}(v)$, is the set of the $4\lceil\alpha \log(n)\sqrt{n}\rceil$ closest nodes to v , for a large enough constant $\alpha > 0$, where ties are broken using the order of node identities. By this definition, if $u \in \text{ball}(v)$ and w is on a shortest path between v and u , then $u \in \text{ball}(w)$.

4.1.1 Color-Sets

In [2], the nodes are partitioned into sets $C_1, \dots, C_{\sqrt{n}}$, called *color-sets*, and, for $i = 1, \dots, \sqrt{n}$, the nodes in color-set C_i are assigned the same color i . For every node $v \in V$, its vicinity ball, $\text{ball}(v)$, contains at least one node from each color-set. To get this, the color of a given node v is determined by a hash function *color* which, given the identity of a node, maps that identity to a color in $\{1, \dots, \sqrt{n}\}$. This mapping from identities to colors is balanced in the sense that at most $O(\log n \sqrt{n})$ nodes map to the same color. A color is chosen arbitrarily, and all nodes with that color are considered to be the landmarks. Let L be the set of landmarks. It holds that $|L| = O(\log n \sqrt{n})$. Also, each node $v \in V$ has at least one landmark in its vicinity ball. We fix, for each vicinity ball $\text{ball}(v)$, an arbitrary landmark, denoted by l_v .

4.1.2 Routing in trees

This construction in [2] makes use of routing schemes in trees. More precisely, they use the results from [8, 19], which states that there exists a shortest-path (name-dependent) routing scheme for trees using names and tables both on $O(\log^2(n)/\log \log(n))$ bits in n -node trees. For a tree T containing node v , let $\text{table}_T(v)$ and $\text{name}_T(v)$ denote the routing table of node v in T , and the name of v in T , as assigned by the scheme in [8].

4.1.3 The routing tables

For any node v , let $T(v)$ be a shortest-path spanning tree rooted at v . Let $P(v, w, u)$ be a path between v and u composed of a shortest path between v and w and a shortest path between w and u . Such a path is said to be *good* for (v, u) if $v \in \text{ball}(w)$, and there exists an edge $\{x, y\}$ along a shortest path between w and u with $x \in \text{ball}(w)$ and $y \in \text{ball}(u)$;

Every node $v \in V$ stores the following information in its routing table $\text{table}(v)$:

- the hash function *color* that maps identities in colors;
- the identity of every node $u \in \text{ball}(v)$, and the port number $\text{next}(v, u)$;
- for every landmark $l \in L$, the routing table $\text{table}_{T(l)}(v)$ for routing in $T(l)$;
- for every node $u \in \text{ball}(v)$, the routing table $\text{table}_{T(u)}(v)$ for routing in $T(u)$;
- the identities of all nodes with same color as v , and, for each such node u , the following additional information:
 - if there are no good paths $P(v, w, u)$, then v stores

$$(\text{name}_{T(l_u)}(l_u), \text{name}_{T(l_u)}(u)). \quad (2)$$

- if there exists a good path $P(v, w, u)$, then let us pick a good path P of minimum length among all good paths; then, let us compare its length $|P|$ with the length $|Q|$ of the path Q composed of a shortest path between v and l_u , and a shortest path between l_u and u ; provide v with

$$(\text{name}_{T(l_u)}(l_u), \text{name}_{T(l_u)}(u))$$

if $|Q| \leq |P|$, and with

$$(\text{name}_{T(v)}(w), x, \text{port}_x(\{x, y\}), \text{name}_{T(y)}(u)) \quad (3)$$

otherwise.

Storing the hash function *color* requires $O(\sqrt{n})$ bits. L and $\text{ball}(v)$ are both of size $\tilde{O}(\sqrt{n})$ bits. Moreover, the number of nodes with identical color is $\tilde{O}(\sqrt{n})$, for every color. Finally, shortest-path routing in any tree can be achieved using tables and names of size $O(\log^2 n / \log \log n)$ bits [8, 19]. It follows that $|\text{table}(v)| = \tilde{O}(\sqrt{n})$, as desired.

4.1.4 Routing

Routing from a source s to a target t is achieved in the following way (see [2] for more details). If $t \in \text{ball}(s)$, or $t \in L$, or s and t have the same color, then s routes to t using the information available in its table. More specifically, if $t \in \text{ball}(s)$ or $t \in L$, then s sets the header of the message as just the identity of the target t . Instead, if s and t have the same color but $t \notin \text{ball}(s) \cup L$, the source s sets the header as one of the two possible cases presented in Eq. (2) and (3). Otherwise, that is, if $t \notin \text{ball}(s) \cup L$ and $\text{color}(s) \neq \text{color}(t)$, node s routes the message towards some node $w \in \text{ball}(s)$ sharing the same color as t . (The color of t can be obtained by hashing the identity of t). The header is set to t , and w will change the header upon reception of this message, according to the rules previously specified. It is shown in [2] that this routing guarantees a stretch 3.

4.2 On the difficulty of locally verifying the scheme in [2]

We note that the routing scheme in [2], as sketched in the previous subsection, has some “*global*” features that makes it plausibly difficult to locally verify, by adding certificates of size $\tilde{O}(\sqrt{n})$ bits at each node. In this subsection, we mention one of these global features, illustrated in the example depicted on Figures 1 and 2.

In both Figures 1 and 2, we are considering routing from a source $s \notin L$ to a target $t \notin L$, of different colors, with $s \notin \text{ball}(t)$ and $t \notin \text{ball}(s)$. Let us assume that node s has color red, while node t has color blue. According to the routing scheme in [2], node s first routes the message towards some blue node $u \in \text{ball}(s)$ to get information about the blue target t . In the example of Fig. 1, node u stores

$$(\text{name}_{T(u)}(s), x, \text{port}(x, y), \text{name}_{T(y)}(t)),$$

in order to guarantee a stretch 3. Instead, in the example of Fig. 2, node u stores

$$(\text{name}_{T(l_t)}(l_t), \text{name}_{T(l_t)}(t)).$$

to guarantee such a small stretch. Verifying locally whether there exists a good path between s and t , which is the condition leading to distinguishing the case where the content of Eq. (2) in Section 4.1.3 must be placed in the table, from the case where the content of Eq. (3) must be

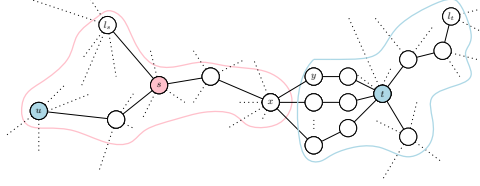


Figure 1: All the nodes in a shortest path between s and t belong to $\text{ball}(s) \cup \text{ball}(t)$.

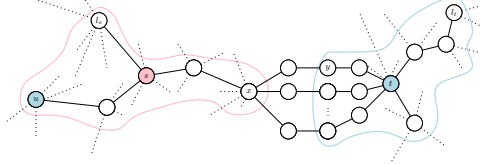


Figure 2: There are nodes in a shortest path between s and t that do not belong to $\text{ball}(s) \cup \text{ball}(t)$.

placed in the table, appears to be a very difficult matter when restricted to certificates of size $\tilde{O}(\sqrt{n})$.

We point out that ignoring the two cases, and systematically storing the content of Eq. (2) would result in a routing scheme that may be locally verifiable. However, its stretch is at least 7. To see why, let us consider the example displayed in Figure 3 where $s \in \text{ball}(t)$ but $t \notin \text{ball}(s)$. The radius of $\text{ball}(s)$ is 50. Let $u \in \text{ball}(s)$ be one of the farthest nodes to s , i.e., $\delta(s, u) = 50$. The radius of $\text{ball}(t)$ is 100 and $\delta(t, l_t) = 100$. Although $\delta(u, t) = 100$, we assume $u \notin \text{ball}(t)$ due to lexicographical order priorities. Finally, assume that $\delta(s, t) = 50$. The worst case route from s to t would then be $s \rightsquigarrow u \rightsquigarrow s \rightsquigarrow t \rightsquigarrow l_t \rightsquigarrow t$. This path is of length

$$\begin{aligned} & \delta(s, u) + \delta(u, s) + \delta(s, t) + \delta(t, l_t) + \delta(l_t, t) \\ &= 2 \times 50 + 50 + 2 \times 100 = 350 \end{aligned}$$

which is $7\delta(s, t)$.

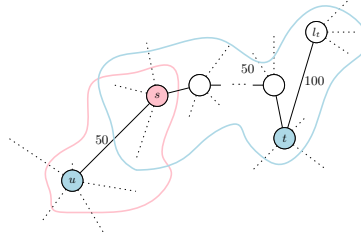


Figure 3: A route of stretch 7.

5 A New Name-Independent Routing Scheme

In this section, we describe and analyze a new name-independent routing scheme, denoted by \mathcal{R} . The section is entirely dedicated to the proof of the following theorem.

Theorem 4 *The name-independent routing scheme \mathcal{R} uses routing tables of size $\tilde{O}(\sqrt{n})$ bits at each node. It guarantees stretch 5, and, using handshaking, its stretch can be reduced to 3. In both cases, \mathcal{R} can be locally verified using certificates of size $\tilde{O}(\sqrt{n})$ bits, using a 1-sided error verification algorithm which guarantees that incorrect tables are detected with high probability⁴.*

5.1 The new routing scheme \mathcal{R}

Our new routing scheme \mathcal{R} borrows ingredients from both [2] and [19]. In particular, the landmarks are chosen as for the (name-dependent) routing scheme in [19], i.e., in such a way that every node v has a cluster $\text{cluster}(v)$ of size at most $4\sqrt{n}$, and the landmarks form a set L of size at most $2 \log n \sqrt{n}$. However, we reinforce the way the nearest landmark to v is selected, by picking the nearest landmark l_v such that

$$l_v = \operatorname{argmin}_{l \in L} \delta(v, l)$$

where ties are broken by choosing the landmark with smallest identity. Also, we slightly reinforce the definition of next : For every two nodes v and t , we set

$$\text{next}(v, t)$$

as the *smallest* port number at v of an edge incident to v on a shortest path between v and t . These two reinforcements of the definitions of landmarks and next guarantee the following.

Lemma 5 *Let $v \in V$, and let l_v be its landmark. Let u be a node on a shortest path between v and l_v . Then $l_u = l_v$.*

Proof. By the choice of the landmarks, if $l_u \neq l_v$ then $\delta(u, l_u) < \delta(u, l_v)$. However, $\delta(v, l_u) \leq \delta(v, u) + \delta(u, l_u)$. So, if $l_u \neq l_v$ then $\delta(v, l_u) < \delta(v, u) + \delta(u, l_v) = \delta(v, l_v)$, contradicting the choice of l_v as the landmark closest to v . \square

Lemma 6 *Let $v \in V$, let l_v be its landmark, and let w be the neighbor of l_v such that $\text{next}(l_v, v) = \text{port}_{l_v}(\{l_v, w\})$. Let u be a node on a shortest path between v and w . We have $\text{next}(l_v, u) = \text{next}(l_v, v)$.*

Proof. By Lemma 5, we have $l_u = l_v$. If $\text{next}(l_v, u) \neq \text{next}(l_v, v)$ there exists a shortest path P from l_v to u that does not go to w but through another neighbor w' of l_v with $\text{next}(l_v, u) = \text{port}_{l_v}(\{l_v, w'\})$. On the other hand, there is a shortest path between l_v and u going through w . So, $\text{port}_{l_v}(\{l_v, w'\}) < \text{port}_{l_v}(\{l_v, w\})$. But the path consisting of P and the shortest path between u and v is also a shortest path between l_v and v . Therefore, by our reinforcement of the definition of next , we get $\text{port}_{l_v}(\{l_v, w'\}) > \text{port}_{l_v}(\{l_v, w\})$, contradiction. \square

⁴Given a family of events $(\mathcal{E}_n)_{n \geq 1}$, event \mathcal{E}_n holds with high probability if $\Pr[\mathcal{E}_n] = 1 - O(1/n^c)$ for $c \geq 1$.

As in [2], each node that is not a landmark is given a color in $\{1, \dots, \sqrt{n}\}$ determined by a hash function, $color$, that maps identities to colors, where at most $O(\log n \sqrt{n})$ nodes map to the same color. Also, each node v has a vicinity ball, $\mathbf{ball}(v)$, that contains the $O(\log(n)\sqrt{n})$ nodes closest to v (breaking ties using identities). This guarantees that, with high probability, for every node v , there is at least one node of each color in $\mathbf{ball}(v)$.

For each color c , where $1 \leq c \leq \sqrt{n}$, we define a set

$$\mathbf{Dir}_c = \{(v, l_v, \mathbf{next}(l_v, v)) : color(v) = c\}$$

which includes the direction to take at l_v for reaching node v along a shortest path, for v of color c .

5.1.1 The routing tables

Every node $v \in V$ then stores the following information in $\mathbf{table}(v)$:

- the hash function $color$ that maps identities in colors;
- the identity of every landmark $l \in L$, and the corresponding port $\mathbf{next}(v, l)$;
- the identity of every node $u \in \mathbf{cluster}(v)$, and the corresponding port $\mathbf{next}(v, u)$;
- the identity of every node $u \in \mathbf{ball}(v)$, and the corresponding port $\mathbf{next}(v, u)$;
- the set $\mathbf{Dir}_{color(v)}$.

Storing the hash function $color$ requires $O(\sqrt{n})$ bits, as we use the same function as in [2]. Since L , $\mathbf{cluster}(v)$, $\mathbf{ball}(v)$, and $\mathbf{Dir}_{color(v)}$ are all of size $\tilde{O}(\sqrt{n})$, we get that $|\mathbf{table}(v)| = \tilde{O}(\sqrt{n})$ as desired.

5.1.2 Routing

Let us consider routing towards a target node t , and let v be the current node. If $t \in \mathbf{cluster}(v)$, then routing to t is achieved using $\mathbf{next}(v, t)$. Notice that, by Lemma 3, routing to t will actually be achieved along a shortest path. Similarly, if $t \in \mathbf{ball}(v)$, then routing to t is achieved using $\mathbf{next}(v, t)$ along a shortest path, and this also holds if t is a landmark. In general, i.e., if t is neither a landmark nor a node of $\mathbf{cluster}(v) \cup \mathbf{ball}(v)$, then node v computes $color(t)$ by hashing the identity of t .

If $color(t) = color(v)$, then v forwards the message towards l_t using the information in $\mathbf{Dir}_{color(v)}$, and including $(l_t, \mathbf{next}(l_t, t))$ in the header of the message so that intermediate nodes carry on routing this message to l_t . At l_t , the message will be routed to t using the information $\mathbf{next}(l_t, t)$ available in the header, reaching a node u_t such that $t \in \mathbf{cluster}(u_t)$. At this point, routing proceeds to t along a shortest path.

If $color(t) \neq color(v)$, then the message is forwarded to an arbitrary node $w \in \mathbf{ball}(v)$ having the same color as t (we know that such a node exists), with w in the header of the message. The message then reach w along a shortest path. At w , we have $color(t) = color(w)$, and thus routing proceeds as in the previous case.

Handshaking. The routing with handshaking to node t proceeds as follows. If $t \in L \cup \mathbf{cluster}(v) \cup \mathbf{ball}(v)$, or $color(v) = color(t)$, then routing proceeds as above. Otherwise, v performs a handshake with a node $w \in \mathbf{ball}(v)$ with $color(w) = color(t)$ in order to get the identity of l_t as well as $\mathbf{next}(l_t, t)$. Then v routes the message to l_t , where it is forwarded to a node u_t such that $t \in \mathbf{cluster}(u_t)$. At this point, routing proceeds to t along a shortest path.

5.2 Stretch of the new routing scheme \mathcal{R}

Let $s, t \in V$ be two arbitrary nodes of the graph. We show that the routing scheme \mathcal{R} routes messages from s to t along a route of length at most $5\delta(s, t)$ in general, and along a route of length at most $3\delta(s, t)$ whenever using handshaking.

As we already observed, if $t \in L \cup \text{cluster}(v) \cup \text{ball}(v)$, then the message is routed to t along a shortest path, i.e., with stretch 1. Otherwise, we consider separately whether the color of s is the same as the color of t , or not.

Assume first that $\text{color}(t) = \text{color}(s)$. Then the message is routed towards l_t along a shortest path, then from l_t to t along a shortest path. The length ℓ of this route satisfies

$$\ell = \delta(s, l_t) + \delta(l_t, t).$$

By the triangle inequality, we get that

$$\ell \leq \delta(s, t) + 2\delta(l_t, t).$$

Since $t \notin \text{cluster}(s)$, we get $\delta(s, t) \geq \delta(t, l_t)$. Therefore

$$\ell \leq 3\delta(s, t).$$

We are left with the case where $\text{color}(s) \neq \text{color}(t)$. Observe first that, with handshaking, the route from s to t will be exactly as the one described in the case $\text{color}(s) = \text{color}(t)$, resulting in a stretch 3. This completes the proof that \mathcal{R} achieves a stretch 3 with handshaking. Without handshaking, the message is forwarded along a shortest path to an arbitrary node $w \in \text{ball}(v)$ having the same color as t , then from w to l_t along a shortest path, and finally from l_t to t along a shortest path. The length ℓ of this route satisfies

$$\begin{aligned} \ell &= \delta(s, w) + \delta(w, l_t) + \delta(l_t, t) \\ &\leq \delta(s, w) + \delta(w, s) + \delta(s, l_t) + \delta(l_t, t) \\ &\leq 2\delta(s, w) + 3\delta(s, t) \\ &\leq 5\delta(s, t) \end{aligned}$$

as desired.

5.3 Local Verification of \mathcal{R}

We show how to verify \mathcal{R} with a verification algorithm `VERIF` using certificates on $\tilde{O}(\sqrt{n})$ bits. Let us define the certificates given to nodes when the routing tables are correctly set as specified by \mathcal{R} .

For each color c , $1 \leq c \leq \sqrt{n}$, let B_c be the number of bits for encoding the set Dir_c , and let $r = \max_{1 \leq c \leq \sqrt{n}} B_c$. We have $r = \tilde{O}(\sqrt{n})$. Let f_1, \dots, f_k be $k = \Theta(\log n)$ hash functions, where each one is mapping sequences of at most r bits onto a single bit. More specifically, each function f_i , $1 \leq i \leq k$, is described as a sequence $f_{i,1}, \dots, f_{i,r}$ of r bits. Given a sequence $D = (d_1, \dots, d_\ell)$ of $\ell \leq r$ bits, we set

$$f_i(D) = \left(\sum_{j=1}^{\ell} f_{i,j} d_j \right) \bmod 2.$$

Hence, if the r bits describing f_i are chosen independently uniformly at random, then, for every two ℓ -bit sets D and D' , we have [21]:

$$D \neq D' \Rightarrow \Pr[f_i(D) = f_i(D')] = \frac{1}{2}$$

Therefore,

$$D \neq D' \Rightarrow \Pr[(f_i(D) = f_i(D'), i = 1, \dots, k)] = 1/2^k.$$

That is, if $k = \beta \lceil \log_2 n \rceil$ with $\beta > 1$, applying the functions f_1, \dots, f_k to both sets D and D' enables to detect that they are distinct, with high probability $1 - 1/n^\beta$.

Each node $v \in V$ stores the certificate composed of the following fields:

- the distances $\{\delta(v, l) : l \in L\}$;
- the distances $\{\delta(v, u) : u \in \text{ball}(v)\}$;
- the set $\{(\delta(v, u), l_u, \delta(u, l_u)) : u \in \text{cluster}(v)\}$;
- the set of k hash functions f_1, \dots, f_k ;
- the set $\{f_i(\text{Dir}_c) : 1 \leq i \leq k, 1 \leq c \leq \sqrt{n}\}$.

The first three entries are clearly on $\tilde{O}(\sqrt{n})$ bits, because of the sizes of L , $\text{cluster}(v)$, and $\text{ball}(v)$. Each function f_i is described by a sequence of $\tilde{O}(\sqrt{n})$ random bits.

The verification algorithm VERIF then proceeds as follows. In a way similar to the proof of Theorem 2, we denote by $H(v)$, $L(v)$, $C(v)$, $B(v)$, $D(v)$, and $\{N(v, t) : t \in L(v) \cup C(v) \cup B(v)\}$ the content of the routing table of v . These entries are supposed to be the hash function *color*, the set of landmarks, the cluster of v , the ball of v , the set $\text{Dir}_{\text{color}(v)}$, and the set of *next-pointers* given to v , respectively. We also denote by d the distance given in the certificates. That is, node v is given a set $\{d(v, t) : t \in L(v) \cup B(v)\}$, and a set $\{(d(v, t), l_t, d(t, l_t)) : t \in C(v)\}$ where l_t is supposed to be the node in $L(v)$ closest to $t \in C(v)$. We also denote by F_1^v, \dots, F_k^v the hash functions given to v in its certificate, and by

$$F(v) = \{F_{i,c}^v : 1 \leq i \leq k, 1 \leq c \leq \sqrt{n}\}$$

the set of $O(\sqrt{n} \log n)$ hash values in the certificates. Of course, if a node does not have a table and a certificate of these forms, then it outputs *reject*. So, we assume that all tables and certificates have the above formats.

Clusters, balls and landmarks (including distances, ports, sizes, etc.) are checked exactly as in Section 3 for the routing scheme in [19]. So, in particular, ignoring the colors and ignoring the minimality of the landmarks' identities, we can assume that, for every node v ,

$$L(v) = L, C(v) = \text{cluster}(v) \text{ and } B(v) = \text{ball}(v),$$

and, for every node $u \in L \cup \text{cluster}(v) \cup \text{ball}(v)$, we have $d(v, u) = \delta(v, u)$, and $N(v, u) = \text{next}(v, u)$ ignoring the minimality of that port number. To check the remaining entries in the routing tables (as well as the previously ignored colors and minimality criteria for the landmarks and the *next-pointers*), VERIF performs the following sequence of steps. At every step, if the test is not passed at some node, then VERIF outputs *reject* at this node, and stops. Otherwise, it goes to the next step. If all tests are passed at a node, that node outputs *accept*. Node v checks that:

1. l_v has the smallest identity among all landmarks closest to v ;
2. it has the same hash function $H(v)$ as its neighbors, that it has one node of each color in $B(v)$, that v appears in $D(v)$, and that all nodes appearing in $D(v)$ have the same color as v ;

3. there exists at least one neighbor u on a shortest path between v and l_v with $N(l_v, v) = N(l_v, u)$; for every neighbor u on the shortest path between v and l_v (implying $l_u = l_v$) with $u \neq l_v$, $N(l_v, v) \leq N(l_u, u)$; and if l_v is a neighbor of v on a shortest path between l_v and v , $N(l_v, v) = \text{port}_{l_v}(\{l_v, v\})$;
4. for every $u \in L \cup B(v) \cup C(v)$, the port number $N(v, u)$ is the smallest among all the ports of edges incident to v on a shortest path between v and u ;
5. it has the same hash functions F_1^v, \dots, F_k^v , as its neighbors;
6. $F_i^v(D(v)) = F_{i, H(v)}^v$ for every $i = 1, \dots, k$;
7. the hash value $F_{i,c}^v$, $1 \leq i \leq k$, $1 \leq c \leq \sqrt{n}$, are identical to the ones of their neighbors.

If v passes all the above tests, then v outputs *accept*, else it outputs *reject*.

We now establish the correctness of this local verification algorithm, that is, we show that it satisfies the specification stated in Section 2.2. First, by construction, if all tables are set according to the specification of \mathcal{R} , then every node running the verification algorithm with the appropriate certificate will face no inconsistencies with its neighbors, i.e., all the above tests are passed, leading every node to accept, as desired. So, it remains to show that if some tables are not what they should be according to \mathcal{R} , then, no matter the certificates assigned to the nodes, at least one node will fail one of the tests with high probability.

If all nodes pass the test of Step 1, then we are guaranteed that not only $L(v) = L$, but also that l_v is indeed the appropriate landmark of v . If all nodes pass the test of Step 2, then it must be the case that $H(v) = \text{color}(v)$, that $B(v) = \text{ball}(v)$ with the desired coloring property, and that

$$D(v) \subseteq \text{Dir}_{\text{color}(v)}.$$

By Lemma 6, we get that if all nodes pass Steps 3 and 4, then $N(v, u) = \text{next}(v, u)$ where the minimality condition is satisfied.

Let us assume that there exists a pair of nodes (u, v) with same color such that

$$(u, l_u, \text{next}(l_u, u)) \in \text{Dir}_{\text{color}(v)} \setminus D(v).$$

By the previous steps, we know that $(u, l_u, \text{next}(l_u, u)) \in D(u)$. Hence, $D(u) \neq D(v)$. On the other hand, if Step 5 is passed by all nodes, then all nodes agree on a set of $k = \Theta(\log n)$ hash function f_1, \dots, f_k . Therefore, assuming that these functions are set at random, we get that, with high probability, there exists at least one function f_i , $1 \leq i \leq k$, such that $f_i(D(v)) \neq f_i(D(u))$.

If all nodes pass Step 6, then in particular $F_i^v(D(v)) = F_{i,c}^v$, and $F_i^u(D(u)) = F_{i,c}^u$ where $c = \text{color}(v) = \text{color}(u)$. We know that $F_{i,c}^v = f_i(D(v))$ and $F_{i,c}^u = f_i(D(u))$, which implies that $F_{i,c}^v \neq F_{i,c}^u$ with high probability, which will be detected at Step 7 by two neighboring nodes in the network. This completes the proof of Theorem 4.

6 Another name-independent Routing Scheme

We have seen that the stretch-3 name-independent routing scheme from [2] does not seem to be locally checkable with certificates of $\tilde{O}(\sqrt{n})$ bits. Our new name-independent routing scheme \mathcal{R} is locally checkable with certificates of $O(\sqrt{n})$ bits, but it has stretch 5. In this section, we focus on the local checkability of the stretch-5 name-independent routing scheme from [3]. We show that, as for the scheme in [2], the routing scheme in [3] do not appear to be locally checkable with certificates of $\tilde{O}(\sqrt{n})$ bits. Moreover, handshaking, which may allow that scheme to become locally verifiable with small certificates, does not reduce its stretch.

6.1 The Routing Scheme of [3]

The scheme in [3] requires $\tilde{O}(\sqrt{n})$ bits of memory per node. It is based on the notion of landmarks, with a set L of landmarks with size $O(\log n \sqrt{n})$. To each node v is also associated a vicinity ball, $\text{ball}(v)$, consisting of the \sqrt{n} closest nodes to v , breaking ties using the lexicographic order of the identities. Each node has at least one landmark in its vicinity ball.

6.1.1 Block partition

Nodes are partitioned into \sqrt{n} blocks, $S_1, \dots, S_{\sqrt{n}}$, each containing \sqrt{n} nodes. Assuming that the identities are from 0 to $n-1$, the set S_i consists of the nodes with identities in $[(i-1)\sqrt{n}, i\sqrt{n})$. Set of blocks are assigned to the nodes. We denote by A_v the set of blocks assigned to node v . We then denote by R_v the set

$$\{u \in V : \exists i \in \{1, \dots, \sqrt{n}\}, u \in S_i \text{ and } S_i \in A_v\}.$$

It is proved in [3] that there exists an assignment of sets of blocks to nodes such that

- for every $v \in V$, and for every block S_i , there exists $u \in \text{ball}(v)$ such that $S_i \in A_u$;
- for every $v \in V$, $|A_v| = O(\log n)$.

6.1.2 The routing tables

This scheme in [3] also uses routing schemes for trees. Similarly to the previously described schemes, for a landmark $l \in L$, we denote by $T(l)$ a shortest-path tree routed at l spanning all the nodes. Also, $\text{table}_T(v)$ and $\text{name}_T(v)$ denote the routing table and name of node v , respectively, for the tree T . Each node $v \in V$ stores the following information in $\text{table}(v)$:

- the identity of every node $u \in \text{ball}(v)$, together with $\text{next}(v, u)$;
- the identity of every node $l \in L$, together with $\text{next}(v, l)$ and $\text{table}_{T(l)}(v)$;
- for every node $u \in \text{ball}(v)$, the set of indexes i such that $S_i \in A_u$;
- the identity of every node $u \in R_v$ together with the identity of $l_{v,u}^*$, and $\text{name}_{T(l_{v,u}^*)}(u)$ where

$$l_{v,u}^* = \text{argmin}_{l \in L} (\delta(v, l) + \delta(l, u)).$$

6.1.3 Routing

Routing from a node s to a node t is achieved in the following way. If $t \in \text{ball}(s) \cup L \cup R_s$, then s routes to t using the information it has in its table. Otherwise, s routes the message towards a node $w \in \text{ball}(s)$ satisfying that $t \in R_w$. Then w routes to t using the information it has in its table, i.e., via the landmark $l_{w,t}^*$ that minimizes the sum $\delta(w, l) + \delta(l, t)$. It is proved in [3] that this routing scheme guarantees stretch 5.

6.2 On the difficulty of locally verifying the scheme in [3]

As for the routing scheme in [2], we note that the routing scheme in [3] sketched in the previous subsection is plausibly difficult to locally verify by adding certificates of size $\tilde{O}(\sqrt{n})$ bits at each node. Indeed, let us consider two nodes u and v with $u \in R_v$, one difficulty consists in verifying that a given landmark $l_{v,u}^*$ is effectively the one that, among all $l \in L$, minimizes the distance $\delta(v, l) + \delta(l, u)$.

Also, we show that, even using handshaking, the scheme in [3] still achieves stretch at least 5. To see why, let us consider the example depicted in Fig. 4, where the node s aims at routing towards the node t . Node $w \in \text{Ball}(s)$ is the only node that has the routing information about node t , i.e., $t \in R_w$ and, for every $u \in \text{ball}(s) \setminus \{w\}$, we have $t \notin R_u$. Therefore, if the handshaking is performed to a node with information about t in $\text{ball}(s)$, then this handshaking is to be performed between s and w . Let us fix some positive values x and ϵ , with $\epsilon \ll x$.

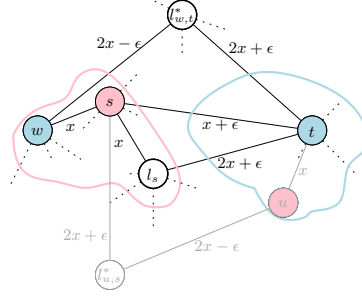


Figure 4: Handshaking does not help.

On Fig. 4, the following holds:

- $t \notin \text{ball}(s)$, and $s \notin \text{ball}(t)$;
- the radius of both $\text{ball}(s)$ and $\text{ball}(t)$ is x ;
- $\delta(s, t) = x + \epsilon$;
- $\delta(s, w) = \delta(s, l_s) = x$ and $\delta(w, l_s) = 2x$;
- $\delta(s, t) = x + \epsilon$;
- $\delta(w, l_{w,t}^*) = 2x - \epsilon$ and $\delta(l_{w,t}^*, t) = 2x + \epsilon$;
- the shortest path between s and $l_{w,t}^*$ goes through w , i.e., $\delta(s, l_{w,t}^*) = \delta(s, w) + \delta(w, l_{w,t}^*) = 3x - \epsilon$.

Note that $\delta(l_s, t) = 2x + \epsilon$, and thus

$$\delta(l_s, t) \leq \delta(l_s, s) + \delta(s, t).$$

Since $l_{w,t}^* = \operatorname{argmin}_{l \in L} (\delta(w, l) + \delta(l, t))$, it follows that

$$\delta(w, l_{w,t}^*) + \delta(l_{w,t}^*, t) \leq \delta(w, l_s) + \delta(l_s, t).$$

In fact, in our example, we have

$$\delta(w, l_{w,t}^*) + \delta(l_{w,t}^*, t) = 4x \leq 4x + \epsilon = \delta(w, l_s) + \delta(l_s, t).$$

Therefore, using handshaking, s queries w that replies with the name of $l_{w,t}^*$. Then s sends the message to t via $l_{w,t}^*$ and the resulting route is $s \rightsquigarrow w \rightsquigarrow l_{w,t}^* \rightsquigarrow t$, which has length:

$$\begin{aligned} \delta(s, w) + \delta(w, l_{w,t}^*) + \delta(l_{w,t}^*, t) \\ = x + 2x - \epsilon + 2x + \epsilon = 5x \end{aligned}$$

which is close to $5\delta(s, t)$ for $\epsilon \ll x$.

The case of handshaking with t also leads to no improvements as far as stretch is concerned. See, e.g., Fig. 4 where t has a red node u in its ball, at distance x from t , and at distance $2x - \epsilon$ from $l_{u,s}^*$ with $l_{u,s}^*$ at distance $2x + \epsilon$ from s . This scenario leads to the very same situation as when the handshaking is performed between s and w .

7 Generalization to larger stretches

In this section, we show that Theorem 2 can be generalized to a whole family of routing schemes achieving a tradeoff between space complexity and stretch. More specifically, Thorup and Zwick describe in [20] an algorithm which, given any integer $k \geq 1$, returns a collection \mathcal{T} of trees such that every node is contained in $O(n^{1/k} \log^{1-1/k} n)$ trees in \mathcal{T} , and, for each pair of nodes, there exists a tree in \mathcal{T} in which these two nodes are connected by a path of length at most $2k - 1$ times the length of a shortest path between them. Thorup and Zwick use this result, combined with the techniques used in their stretch-3 routing scheme, to produce a family of routing schemes indexed by $k \geq 2$. For any $k \geq 2$ the routing scheme in this family uses a memory of $\tilde{O}(n^{1/k})$ bits at each node, and has stretch $4k - 5$ (see [19]). The stretch can be reduced to $2k - 1$ using handshaking.

Theorem 7 *For any $k \geq 2$, the stretch- $(4k-5)$ routing scheme by Thorup and Zwick in [19] can be locally verified using certificates of size $\tilde{O}(n^{1/k})$ bits. This also holds for their handshaking-based stretch- $(2k-1)$ routing scheme.*

We provide below a sketch of proof for Theorem 7. For this purpose, we recall the family of routing schemes in [19].

7.1 The Routing Scheme Family of [19]

The design of this family combines techniques described in [20] and [19].

7.1.1 A hierarchy of landmarks

The design is based on the construction of a hierarchy of landmarks

$$L_0 \supseteq L_1 \supseteq \dots \supseteq L_{k-1} \supseteq L_k$$

where $L_0 = V$, $L_k = \emptyset$, and $|L_i| = \tilde{O}(n^{1-i/k})$ for $i = 0, \dots, k-1$. One way of constructing such set is to place each element of L_{i-1} into L_i independently with probability $n^{-1/k}$. This way, L_{k-1} is of size $O(n^{1/k} \log n)$.

For $0 \leq i \leq k-1$, let $l_i(v)$ be the nearest landmark of node v among all the landmarks in L_i . Notice that, for every node v , we have $l_0(v) = v$. Each node v has a bunch associated to it, defined as follows.

$$\text{bunch}(v) = \bigcup_{0 \leq i < k} \{u \in L_i \setminus L_{i+1} : \delta(v, u) < \delta(v, l_i(v))\}.$$

Note that, since $L_k = \emptyset$, we have $\delta(v, l_k(v)) = +\infty$ for every $v \in V$, and thus, for every node v ,

$$L_{k-1} \subseteq \text{bunch}(v).$$

It can be shown [20] that, for any node v ,

$$|\text{bunch}(v)| = \tilde{O}(n^{1/k}).$$

The aforementioned collection \mathcal{T} of trees is actually composed on n trees, each rooted at a different node. We denote by $T(v)$ the tree in \mathcal{T} that is rooted at v . Moreover, \mathcal{T} satisfies that, for every node v , and every $u \in \text{bunch}(v)$, node v belongs to $T(u)$. The same also holds for the landmarks, that is, for every node v , and every landmark $l_i(v)$, node v belongs to $T(l_i(v))$.

Each node has also a cluster associated to it, defined as follows.

$$\text{cluster}(v) = \bigcup_{0 \leq i < k} \{u \in L_i \setminus L_{i+1} : \delta(v, u) < \delta(u, l_i(u))\}$$

It can be shown [19] that landmarks can be set in such a way that, for every v ,

$$|\text{cluster}(v)| = O(n^{1/k}).$$

7.1.2 Names and tables

Each node v stores in its table, $\text{table}(v)$, the following information:

- the identities of each node $u \in \text{bunch}(v)$;
- for every $u \in \text{bunch}(v)$, the routing table $\text{table}_{T(u)}(v)$ and the name $\text{name}_{T(u)}(v)$ of node v ;
- for every $u \in \text{cluster}(v)$, the name $\text{name}_{T(v)}(u)$;
- the identity of the landmark $l_i(v) \in L_i$, $0 \leq i \leq k-1$;
- the routing table $\text{table}_{T(l_i(v))}(v)$.

All this information can be stored at v using $\tilde{O}(n^{1/k})$ bits. The $O(\log n)$ -bit name of node v is set as

$$\text{name}(v) = \left((l_i(v), \text{name}_{T(l_i(v))}(v)), i = 0, \dots, k-1 \right).$$

7.1.3 Routing

The routing from a source s to a target t proceeds as follows.

If $t \in \text{bunch}(s)$, or $t = l_i(v)$ for some i , $0 \leq i \leq k-1$, then s routes to t along a shortest path, using the information in its table, i.e., using $\text{table}_{T(t)}(s)$. If $t \in \text{cluster}(s)$, then s routes to t along a shortest path using the information in its table, i.e., using $\text{name}_{T(s)}(t)$ on the tree rooted at s . (Note that if $t \in \text{cluster}(s)$, then $s \in \text{bunch}(t)$, and, more generally, $s \in \text{bunch}(u)$

for any u on a shortest path between s and t , from which it follows that u also knows how to route to t using $\text{table}_{T(s)}(u)$.

Otherwise, node s scans all identities $l_i(t)$ in $\text{name}(t)$, and, among all such landmarks in $\text{bunch}(s)$, picks the one with the smallest index i , i.e., picks $l_i(t)$ where

$$i = \min\{j \in \{0, \dots, k-1\} : l_j(t) \in \text{bunch}(s)\}.$$

Then s routes to t using $\text{name}_{T(l_i(t))}(t)$ in the tree $T(l_i(t))$ — this information can be extracted from the name of t . Note that there exists at least one landmark $l_i(t) \in \text{bunch}(s)$ because $L_{k-1} \in \text{bunch}(v)$ for every node v .

It is proved in [19] that this routing scheme guarantees a stretch $4k-5$.

Handshaking. In order to reduce the stretch from $4k-5$ to $2k-1$, the source s performs handshaking by querying node t , providing it with $\text{name}(s)$. The target t computes

$$j = \min\{\ell \in \{0, \dots, k-1\} : l_\ell(s) \in \text{bunch}(t)\}.$$

Then t chooses between $l_i(t)$ and $l_j(s)$ the one that has the minimum index among the two. That is, it picks

$$l^* = \begin{cases} l_i(t) & \text{if } i \leq j \\ l_j(s) & \text{otherwise.} \end{cases}$$

and routing between s and t is performed along the tree $T(l^*)$, guaranteeing stretch $2k-1$.

7.2 Sketch of proof of Theorem 7

It is not hard to check that one can locally verify the $O(\log^2 n / \log \log n)$ -bit routing tables in [8] for routing in trees with certificates of size $O(\log^2 n / \log \log n)$ bits at each node. As for the basic case of stretch-3 routing with tables of size $\tilde{O}(\sqrt{n})$ bits at each node, each node has to check the correctness of its bunch, its cluster, and its set of landmarks. For this purpose, the certificate of every node v stores

- the distance $\delta(v, u)$ for every $u \in \{l_0(v), \dots, l_{k-1}(v)\}$, and every $u \in \text{cluster}(v) \cup \text{bunch}(v)$;
- the distance $\delta(u, l_u)$ for every node $u \in \text{cluster}(v)$.

Each node v then checks that:

1. its bunch, its cluster, and its set of the nearest landmarks are of the appropriate size;
2. the set of landmarks L_{k-1} is the same for every node;
3. $l_i(v)$ is truly the nearest landmark of level i of node v ;
4. all nodes in $\text{bunch}(v)$ satisfy the definition of bunch, and that there are no missing nodes;
5. the nodes in $\text{cluster}(v)$ satisfy the definition of cluster, and that there are no missing nodes.

All these tests can be applied at each node by using the same techniques as to verify the stretch-3 routing scheme of [19], and the correctness follows by the same arguments as those given in Section 3.

8 Conclusion and Further Work

We have shown that it is possible to verify routing schemes based on tables of size $\tilde{O}(\sqrt{n})$ bits using certificates with sizes of the same order of magnitude as the space consumed by the routing tables. The stretch factor is preserved, but to the cost of using handshaking mechanisms for name-independent routing. We do not know whether there exists a stretch-3 name-independent routing scheme, with tables of size $\tilde{O}(\sqrt{n})$ bits, that can be verified using certificates on $\tilde{O}(\sqrt{n})$ bits. Our new routing scheme, which is verifiable with certificates on $\tilde{O}(\sqrt{n})$, has stretch 3 only if using handshaking (otherwise, it has stretch 5). Moreover, the certification of our routing scheme is probabilistic, and it would be of interest to figure out whether deterministic certification exists for some stretch-3 name-independent routing scheme, with tables and certificates on $\tilde{O}(\sqrt{n})$ bits.

Interestingly, our result for stretch-3 name-dependent routing can be extended to larger stretches. Namely, by using the same techniques as for stretch 3, we can show that the family of routing schemes in [19] using, for every $k \geq 1$, tables of size $\tilde{O}(n^{1/k})$ with stretch at most $4k - 5$ (or $2k - 1$ using handshaking) are verifiable with certificates of size $\tilde{O}(n^{1/k})$. However, we do not know whether such a tradeoff between table sizes and stretches can be established for verifiable *name-independent* routing schemes. Specifically, are the existing families of name-independent routing schemes using, for every $k \geq 1$, tables of size $\tilde{O}(n^{1/k})$ with stretch at most $O(k)$, verifiable with certificates of size $\tilde{O}(n^{1/k})$? If not, is it possible to design a new family of verifiable name-independent routing schemes satisfying the same size-stretch tradeoff?

Acknowledgements: The authors are thankful to Cyril Gavoille and Laurent Viennot for discussions about the content of this paper.

References

- [1] I. Abraham, C. Gavoille, D. Malkhi: *On space-stretch trade-offs: upper bounds*. In Proc. 18th ACM Symposium on Parallel Algorithms and Architectures (SPAA), pp217-224, 2006.
- [2] I. Abraham, C. Gavoille, D. Malkhi, N. Nisan, M. Thorup: *Compact name-independent routing with minimum stretch*. In Proc. 16th ACM Symposium on Parallel Algorithms and Architecture (SPAA), pp20-24, 2004.
- [3] M. Arias, L. J. Cowen, K. A. Laing, R. Rajaraman, O. Taka: *Compact routing with name independence*. In Proc. 15th ACM Symposium on Parallel Algorithms and Architectures (SPAA), pp184-192, 2003.
- [4] B. Awebuch, A. Bar-Noy, N. Linial, D. Peleg: *Compact distributed data structures for adaptive routing*. In Proc. 21st ACM Symposium on Theory of Computing (STOC), pp479-489, 1989.
- [5] B. Awebuch, D. Peleg: *Sparse partitions*. In Proc. 31st IEEE Symposium on Foundations of Computer Science (FOCS), pp503-513, 1990.
- [6] S. Chechik: *Compact routing schemes with improved stretch*. In Proc. ACM Symposium on Principles of Distributed Computing (PODC), pp33-41, 2013.
- [7] L. Feuilloley, P. Fraigniaud: *Survey of Distributed Decision*. Bulletin of the EATCS **119**, pp41-65, 2016.
- [8] P. Fraigniaud, C. Gavoille: *Routing in trees*. In Proc. 28th Int. Colloquium on Automata, Languages and Programming (ICALP), pp757-772, 2001.

- [9] P. Fraigniaud, C. Gavoille. *A space lower bound for routing in trees*. In Proc. 19th Symposium on Theoretical Aspects of Computer Science (STACS), Springer LNCS 2285, pp65-75, 2002.
- [10] P. Fraigniaud, A. Korman, D. Peleg: *Towards a complexity theory for local distributed computing*. J. ACM **60**(5): 35:1-35:26 (2013)
- [11] M. Göös, J. Suomela: *Locally Checkable Proofs in Distributed Computing*. Theory of Computing **12**(1): 1-33 (2016)
- [12] C. Gavoille: *Routing in distributed networks: overview and open problems*. ACM SIGACT News **32**(1):36-52, 2001.
- [13] C. Gavoille, M. Gengler: *Space-efficiency for routing schemes of stretch factor three*. Journal of Parallel and Distributed Computing **61**(5):679-687, 2001.
- [14] C. Gavoille, D. Peleg: *Compact and localized distributed data structures*. Distributed Computing **16**(2-3):111-120, 2003.
- [15] C. Gavoille, S. Perennes: *Memory Requirements for Routing in Distributed Networks*. In Proc. 15th ACM Symp. on Principles of Distributed Computing (PODC), pp125-133, 1996.
- [16] A. Korman, S. Kutten, D. Peleg: *Proof labeling schemes*. Dist. Computing **22**(4): 215-233 (2010)
- [17] D. Peleg, E. Upfal: *A Tradeoff between Space and Efficiency for Routing Tables*. In Proc. 20th ACM Symp. on Theory of Comp. (STOC), pp43-52, 1988.
- [18] N. Santoro, R. Khatib: *Labelling and Implicit Routing in Networks*. Comput. J. **28**(1):5-8, 1985.
- [19] M. Thorup, U. Zwick: *Compact routing schemes*. In Proc. 13th ACM Symposium on Parallel Algorithms and Architectures (SPAA), pp1-10, 2001.
- [20] M. Thorup, U. Zwick: *Approximate distance oracles*. In Proc. 33rd ACM Symposium on Theory of Computing (STOC), pp183-192, 2001.
- [21] A.C. Yao: *Some Complexity Questions Related to Distributive Computing*. In Proc. 11th ACM Symposium on Theory of Computing (STOC), pp209-213, 1979.